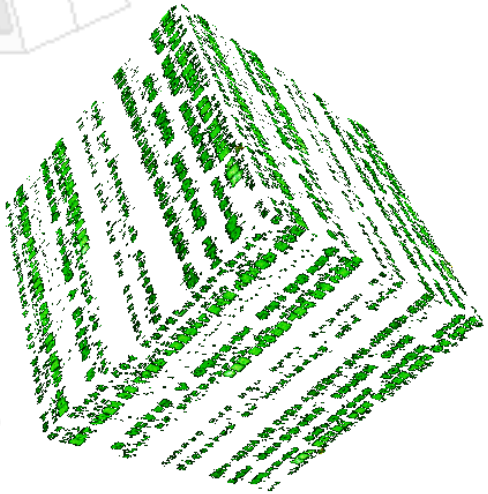

Fitting Arrays into Big Data Analytics

ISO/IEC JTC 1 Study Group on Big Data
2014-mar-19

Dimitar Mišev
rasdaman GmbH | Jacobs University Bremen
misev@rasdaman.com

Array DB R&D

- Large-Scale Scientific Information Systems research group
  - Focus: large-scale n-D raster services & beyond
  - www.jacobs-university.de/lsis
- Main results:
  - Array DBMS, rasdaman
    - Spin-off company: rasdaman GmbH
  - Geo service standards: Chair, OGC raster-relevant working groups, editor of 15 „Big Data“ stds

Who Needs Arrays?

- Sampled, discretized data, eg: Sensor, image, model, & statistics data
  - **Earth:** Geodesy, geology, hydrology, oceanography, climate, earth system, ...
  - **Space:** optical / radio astronomy, cosmological simulation, planetary science, ...
  - **Life:** Pharma/chem, healthcare / bio research, bio statistics, genetics, ...
  - **Engineering & research:** Simulation & experimental data in automotive/shipbuilding/ aerospace industry, turbines, process industry, ...
  - **Management/Controlling:** Decision Support, OLAP, Data Warehousing, census, statistics in industry and public administration, ...
  - **Multimedia:** distance learning, prepress, ...

- „80% of all data have some spatial connotation“ [C&P Hane, 1992]

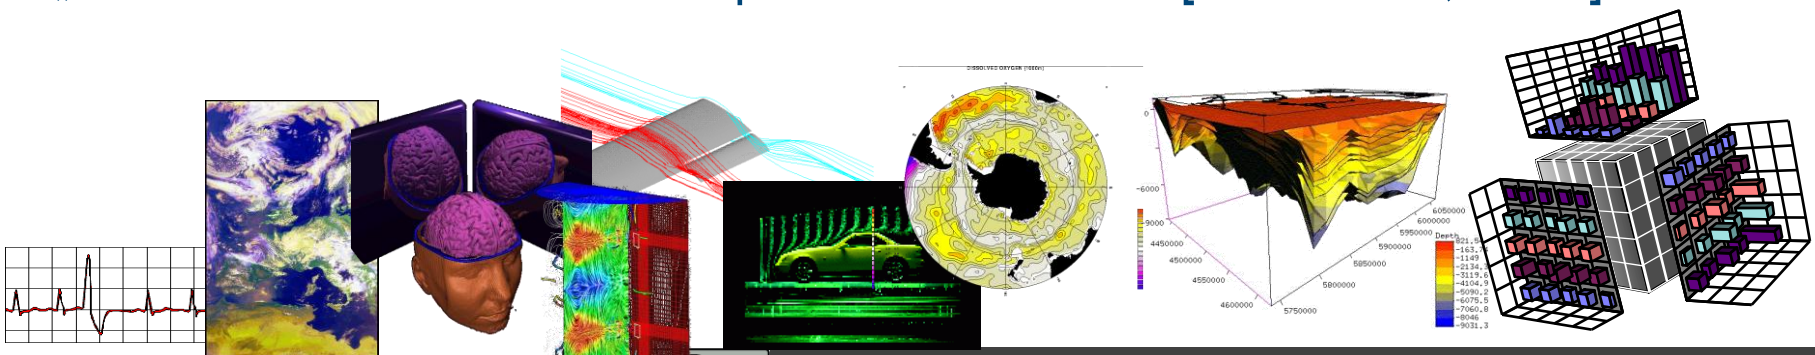

Array Analytics

- **Array Analytics** :=
  Efficient analysis on multi-dimensional arrays of a size several orders of magnitude above the evaluation engine's main memory
- Essential data property: *n*-dimensional Euclidean neighborhood
  - Multidimensional array := $f: E^n \rightarrow V$

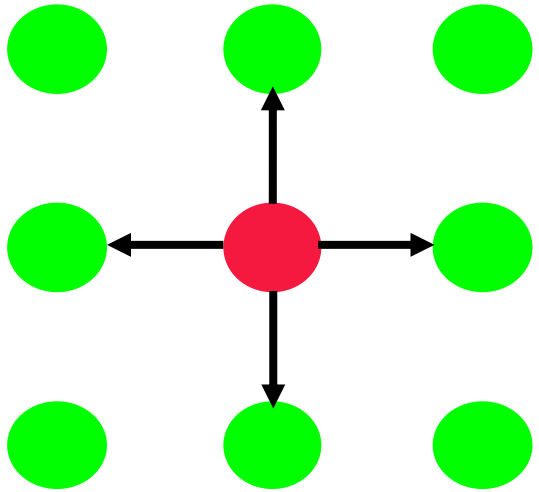

Let's Take a Closer Look...

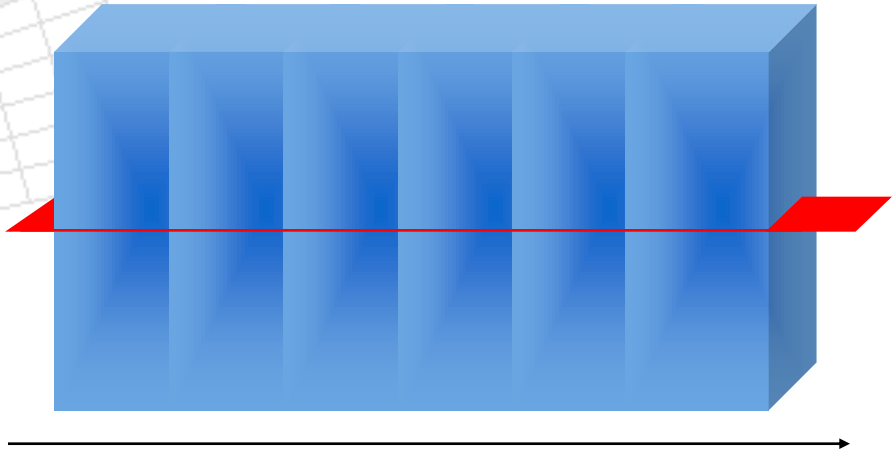

- Divergent access patterns for ingest and retrieval
- Server must mediate between access patterns

rasdaman raster data management

JACOBS UNIVERSITY

Conceptual Modelling of Arrays in Databases

Integrating Arrays Into SQL Traffic

- Goal: minimally invasive integration into SQL
  - SQL'99 onwards supports array attributes (but 1D and no good ops)
- Arrays modeled as attributes (SQL, rasdaman, PostGIS raster, SciSPARQL, ...)

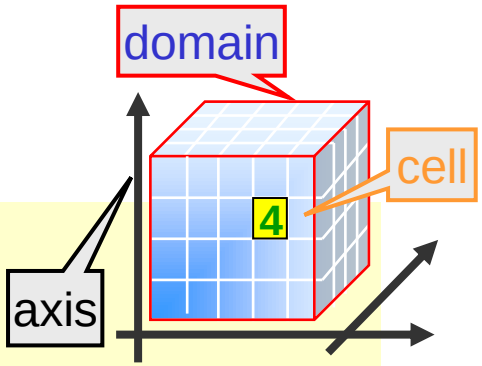


```
create table LandsatScenes(
  id integer not null,
  acquired date,
  scene row( band1 integer, ..., band5 integer )
        array [ x(0:4999), y(0:4999)  ]
)
```

- Arrays modeled as tables (alternative approach: SciQL, SciDB)

```
create array LandsatScene1(
  x int dimension[0:1:4999],
  y int dimension[0:1:4999],
  band1 integer, ..., band5 integer,
)
```

Integrating Arrays Into SQL Traffic / contd.

- selection & section

```
select c.data[ *:*, 100:200, *:*, 42 ]
from   ClimateSimulations as c
```

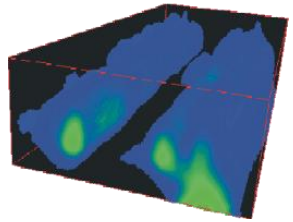

- result processing

```
select scene * (scene.band3 > 130)
from   LandsatScenes as ls
```

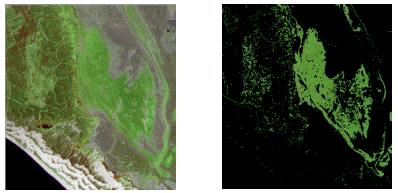

- search & aggregation

```
select mri.data
from   MRI as img, masks as m
where  some_cells( mri.data > 250 and m.valid )
```

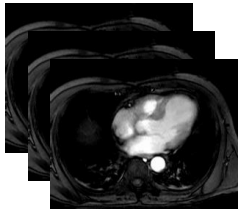

- data format conversion

```
select encode( c.data[*:*,*:*,100,42], „image/png“ )
from   ClimateSimulations as c
```

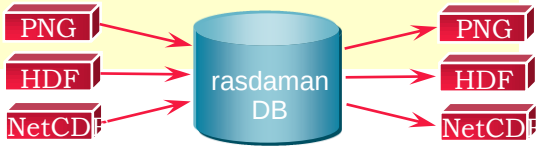

Array Query Operators: Concept

- General array constructor
  - **array** [x(x1:x2), y(y1:y2)]
    values a[x] * b[y]
- Subsetting: trim & slice
  - a[x0:x1, y0:y1, z]
- Constructor shorthands
  - *„for every cell type allowed in arrays, all operations on that type shall be provided on arrays, too“*
  - a + b, log(a), a > b, ...

- General array aggregator
  - **aggregate** max
    over [x(x1:x2), y(y1:y2)]
    where a[x,y] **not null**
    using a[x,y] – a[x-1,y]
- Aggregator shorthands
  - some_cells(b), all_cells(b)
  - count_cells(a), add_cells(a), ...
  - min_cells(a), max_cells(a)

- Array **nest** / **unnest**
  - *for sort, group by, ...*

Architecture I: Storage Mapping

Tiled Array Storage

- partition multidimensional object
  → multidimensional tiles
  - Tile = subarray [Baumann, VLDB Journal 1994]
  - „chunks“ [Sarawagi, DeWitt, ...]
  - SciDB: overlapping borders
  - SciQL: linearized (column store)
  - PostGIS Raster: 2-D tiles

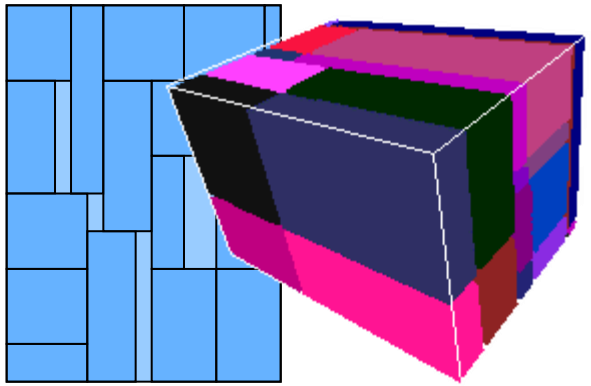

- Tiles = unit of disk access
  - BLOB in RDBMS ...or file
  - Compression, spatial index

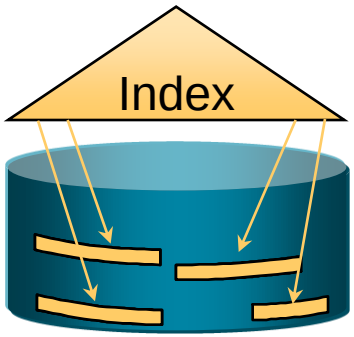

Configurable Tiling

- tiling strategies as service tuning [Furtado]:
  - regular

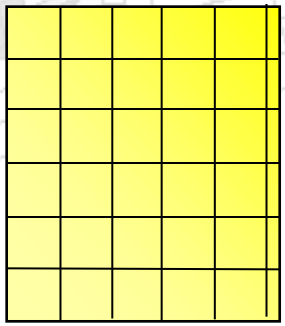

  - directional

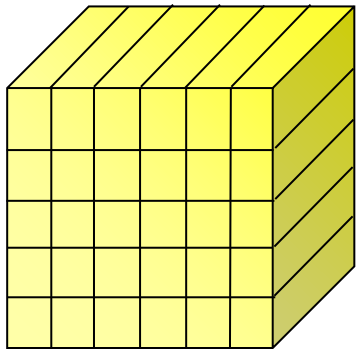

  - area of interest

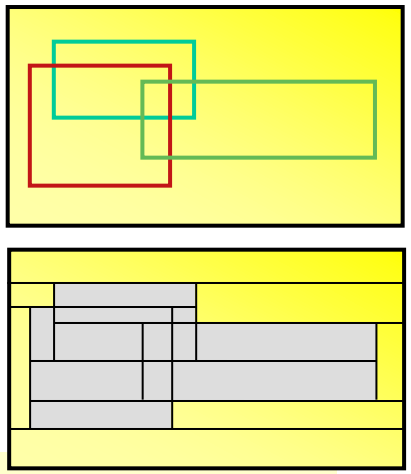

  - rasdaman storage layout language

```
insert into LandsatScenes
  values ...
  tiling area of interest [0:20,0:40], [45:80,80:85]
  tile size 1000000
  index d_index
  storage array compression zlib
```

- SciDB: overlapping tile borders

Architecture II: Query Processing

Query Processing: Overview

- Powerful Array Algebra
- Clear separation:
  set vs array trees
- Extensive optimization
- Tile-based evaluation

```
select a < avg_cells( b + c )
from   a, b, c
```

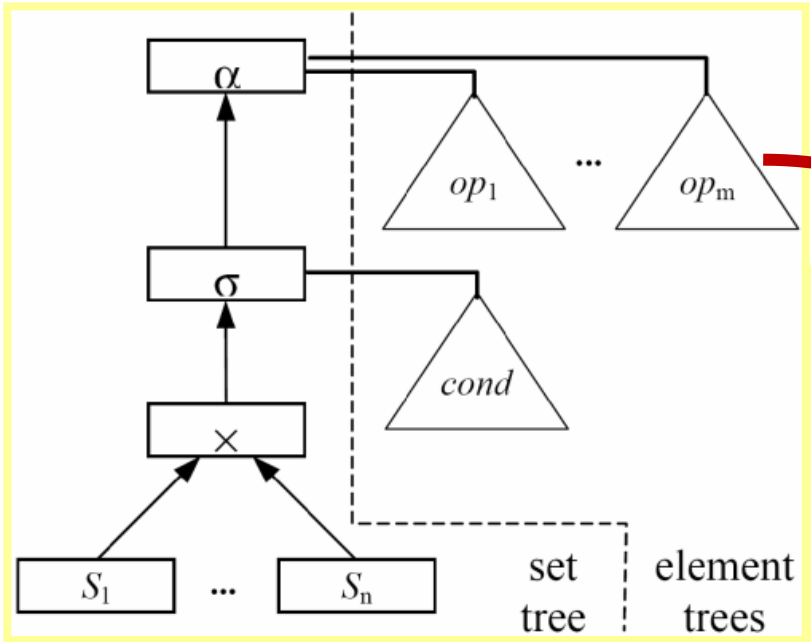

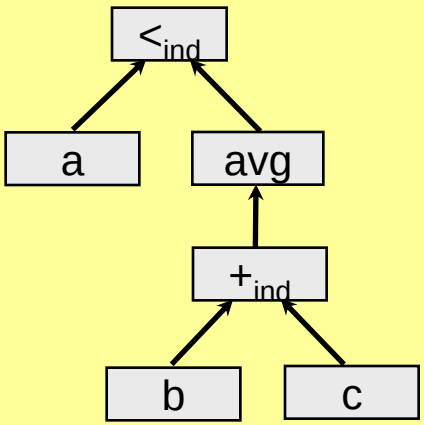

no messing with set tree

Query Rewriting

```
select avg_cells( a + b )
from   a, b
```

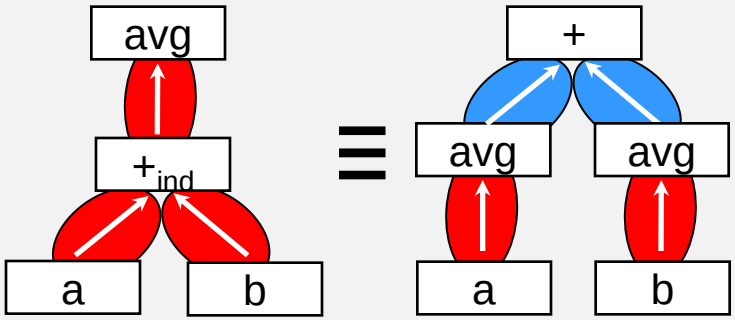

Tile stream
high traffic

Scalar stream
low traffic

```
select avg_cells( a )
     + avg_cells( b )
from   a, b
```

- *understood:*
  heuristic optimization
  – 150 rules in rasdaman [Ritsch 2002]
- *partially understood:*
  cost-based optimization

Heterogeneous Target Code Generation

- Observation: interpreted mode slows down
- Approach:
  - cluster suitable operations
  - compile & dynamically bind
- Benefit:
  - Speed up complex, repeated operations
- Variation:
  - compile code for GPU (later)

```
select x*x*...*x
from float_matrix as x
```

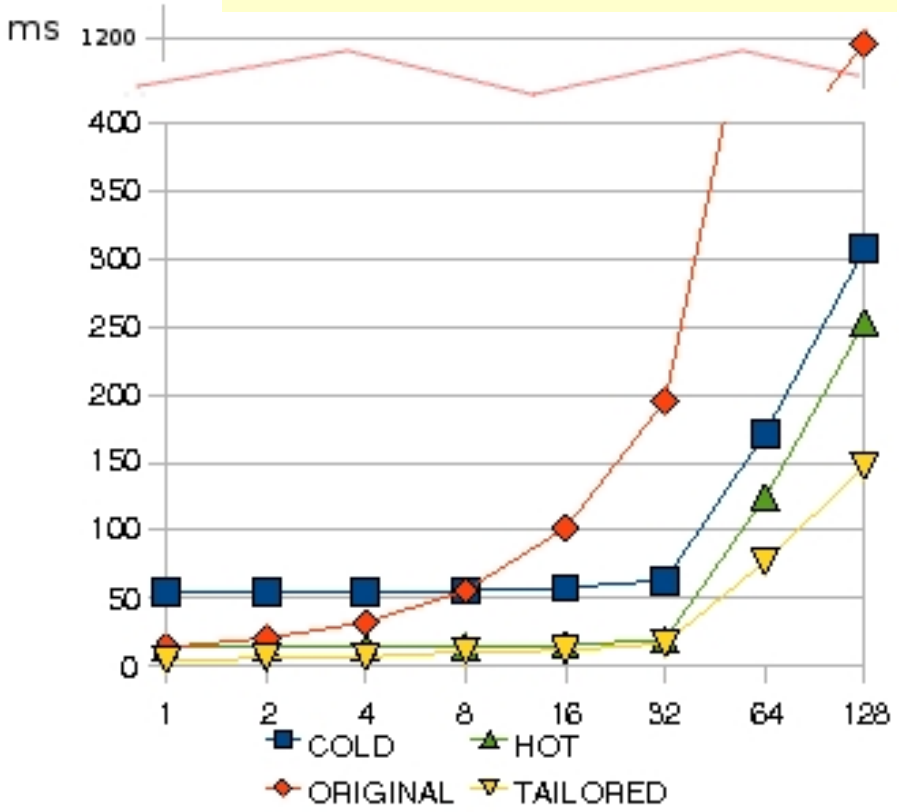

Times [ms] for 512^{2*n} ops

[Jucovschi, Stancu-Mara 2008]

Federated Query Processing

- Heterogeneous rasdaman peer networks
- Manifold criteria possible:
  data location, minimal transport,
  current load, ...

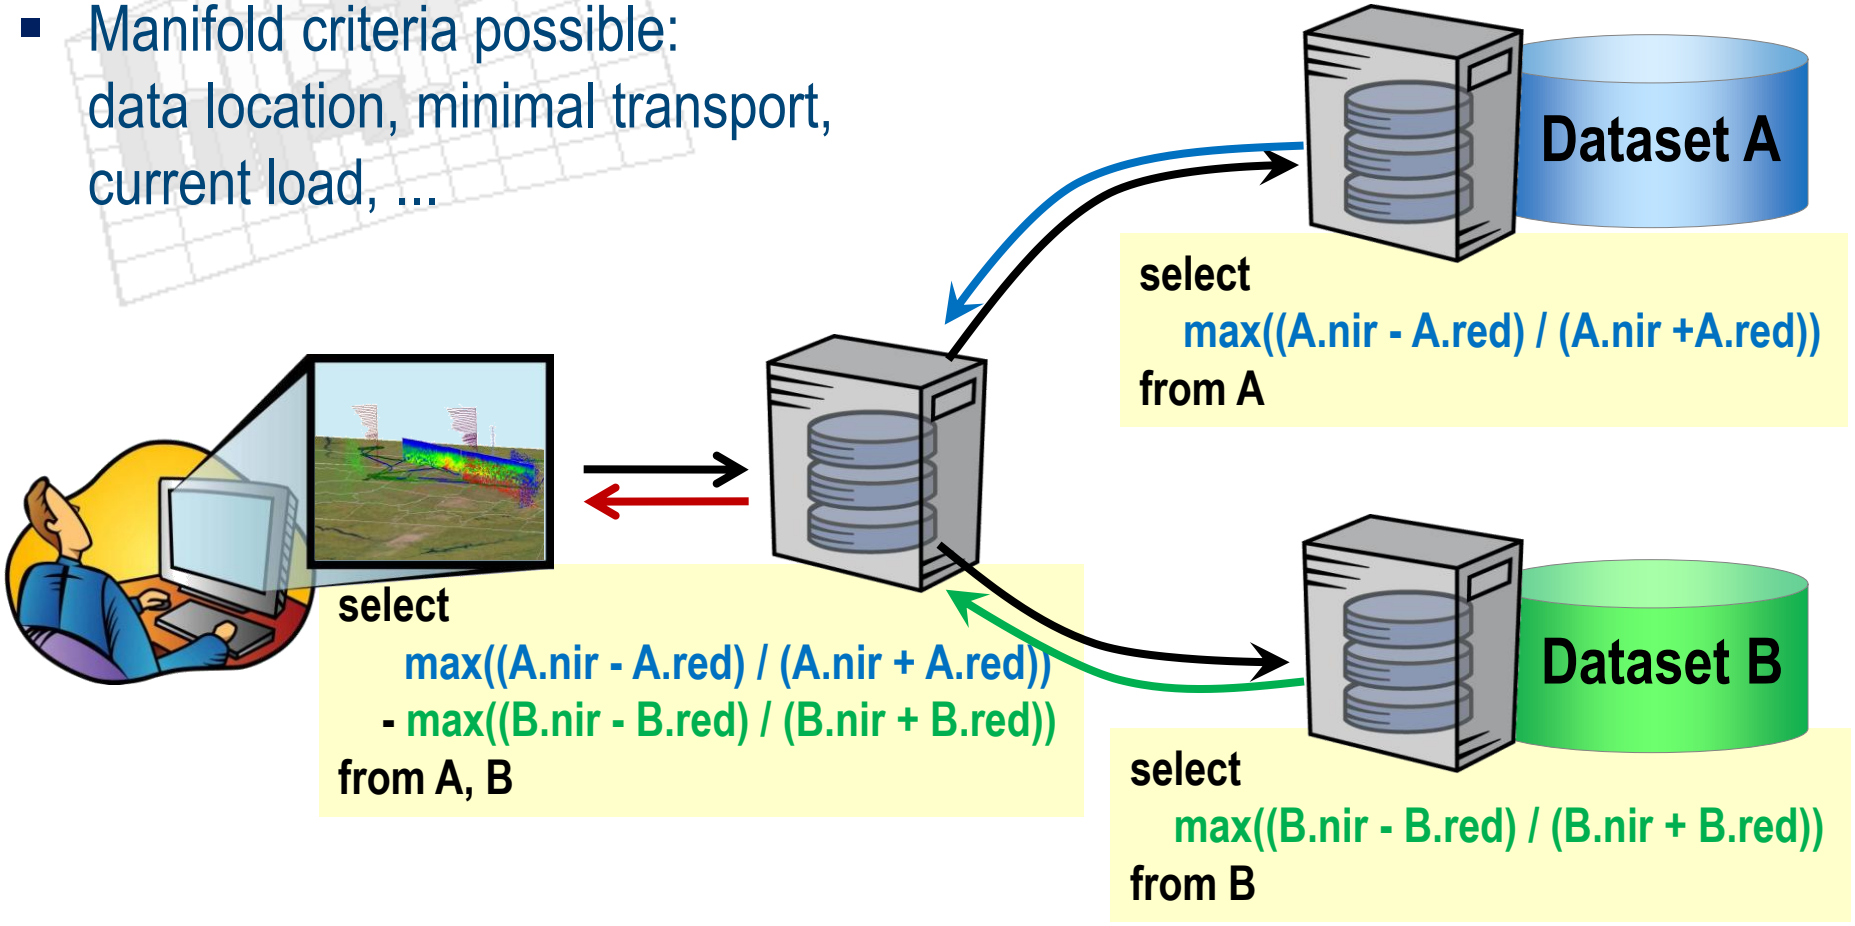

[Owonibi 2012]

rasdaman
raster data management

JACOBS
UNIVERSITY

Array Database Applications

rasdaman raster data management

JACOBS UNIVERSITY

EarthServer: *Big Earth Data Analytics*

- Scalable On-Demand Processing for the Earth Sciences
  - EU FP7-INFRA, 11 partners, Sep 2011 – Aug 2014, 5.85 mEUR
- strictly open standards, centered around OGC coverages
- Big Science Data technology: rasdaman
- Setting up ~300 TB operational Earth & Planetary science data services

Cryospheric Science
landcover mapping

EOX

Airborne Science
high-altitude long-endurance drones

NASA

Atmospheric Science
climate variables

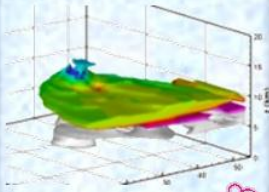

MEEO
Meteorological Environmental Earth Observation

Geology
geological models

British Geological Survey
NATURAL ENVIRONMENT RESEARCH COUNCIL

Oceanography
marine model runs + in-situ data

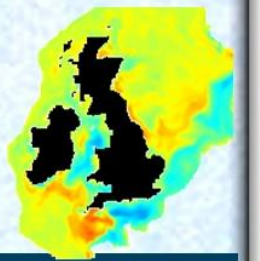

PML PLYMOUTH MARINE LABORATORY

Planetary Science
Mars geology

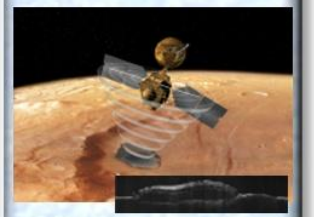

JACOBS UNIVERSITY

On-The-Fly Visualization

[data courtesy BGS, ESA]

- COMETA: mobile WCS client
- JacobsUni: web client toolkit
- Fraunhofer: 3D web client

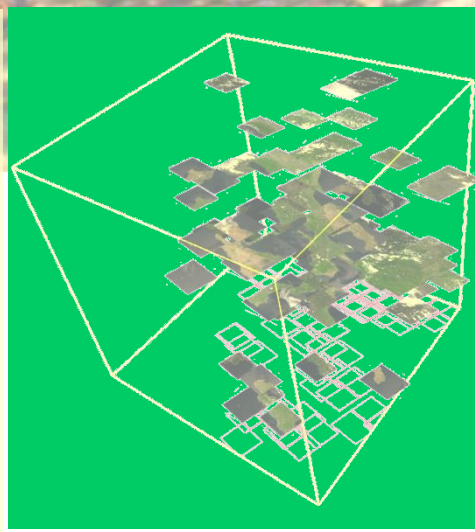

```
select
  encode(
    struct {
      red:   (char) s.b7[x0:x1,x0:x1],
      green: (char) s.b5[x0:x1,x0:x1],
      blue:  (char) s.b0[x0:x1,x0:x1],
      alpha: (char) scale( d, 20 )
    },
    "image/png"
  )
from LandsatScenes as s, DEM as d
```

[JacobsU, Fraunhofer 2012]

OGC Web Coverage Processing Service

- OGC Web Coverage Processing Service (WCPS) - adopted 2008
  = high-level grid coverage filtering & processing language

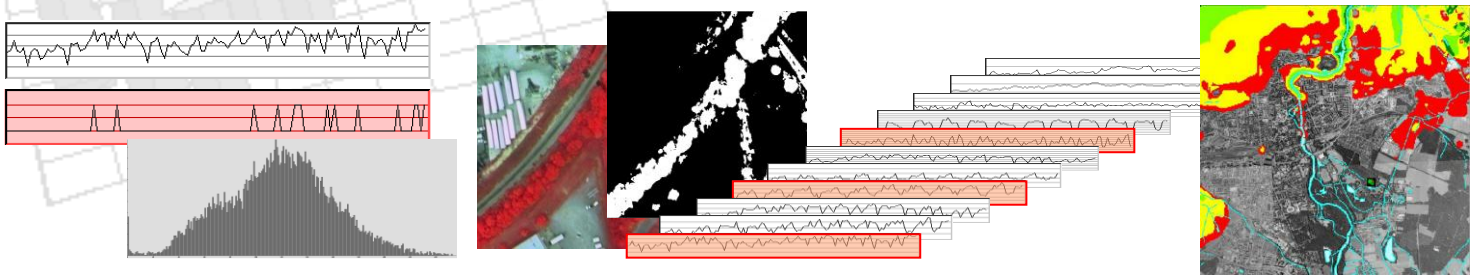

- "From MODIS scenes M1, M2, M3: difference between red & nir, as TIFF"
  - …but only those where nir exceeds 127 somewhere

```
for $c in ( M1, M2, M3 )
where
    some( $c.nir > 127 )
return
    encode(
        $c.red - $c.nir,
        "image/tiff"
    )
```

→ $(\text{tiff}_A, \text{tiff}_C)$

Human Brain Imaging

- Goal: understand structural-functional relations in human brain
- Experiments capture activity patterns (PET, fMRI)
  - Temperature, electrical, oxygen consumption, ...
  - → lots of computations → „activation maps“
- Example: “*a parasagittal view of all scans containing critical Hippocampus activations, TIFF-coded*.“

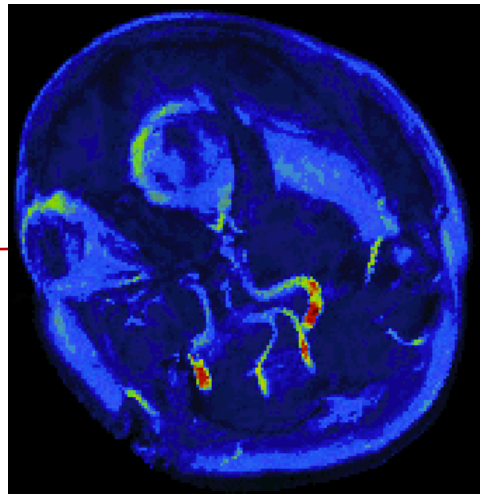

```
select encode(ht[ $1, *:*, *:* ], "tiff")
from   HeadTomograms as ht,
       Hippocampus as mask
where  count_cells( ht > $2 and mask ) /
       count_cells( mask ) > $3
```

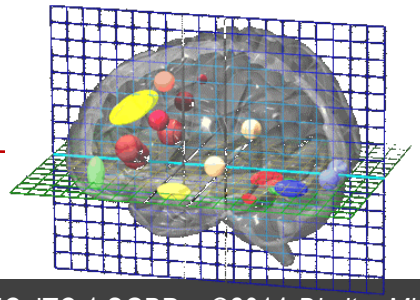

`$1` = slicing position, `$2` = intensity threshold value, `$3` = confidence

On-Board Query Intelligence

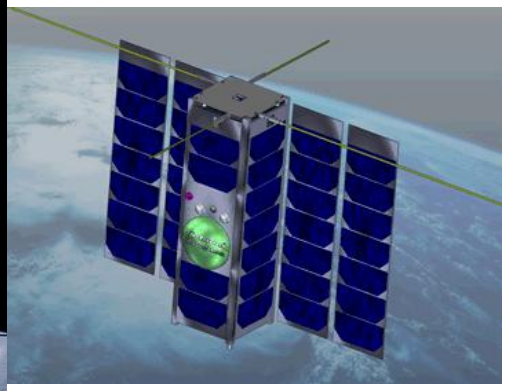

ESA

Democratize direct data access

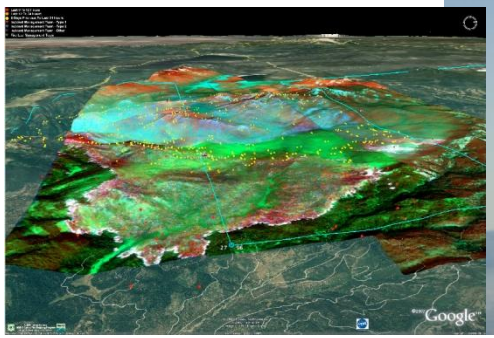

NASA

[imagery courtesy ESA, NASA]

Related Work

Related Work

- Hadoop: generic programming model
  - SciHadoop adds array front-end
  - Requires substantial rewriting of algorithms – needs coding experience

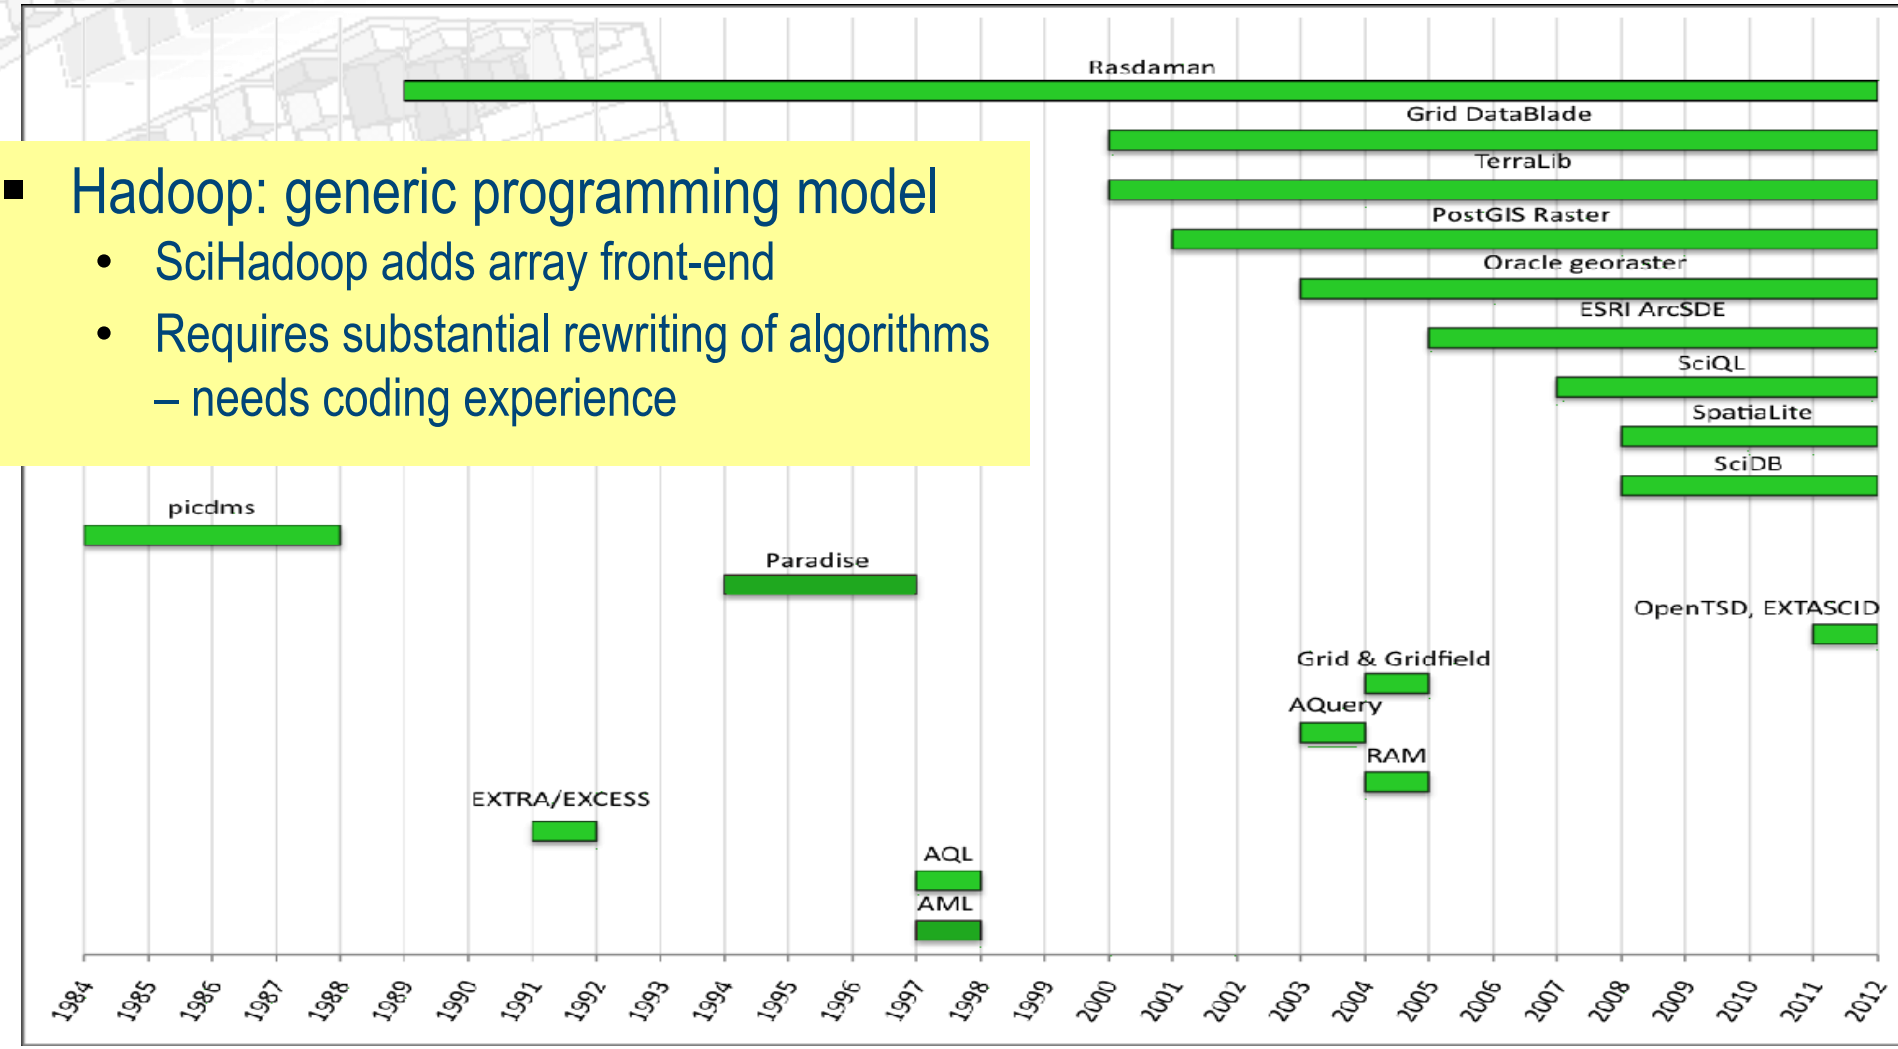

rasdaman raster data management

JACOBS UNIVERSITY

Wrap-Up

Conclusion

- Arrays, graphs major next challenge in conceptual modelling
  - Typically, „Big Data“
  - Found in earth / space / life sciences, business, ...
- Currently databases bypassed, but can contribute substantially
  - Flexibility, scalability, information integration, ...
- Array Databases: adding support for massive n-D arrays
  - Emerging „next wave“ – cf XLDB, Array Databases workshop (www.rasdaman.com/ArrayDatabases_Workshop)
- rasdaman: pioneer Array DBMS
  - Impacting geo & database standards